



Search Results

Search Results for: [((look* or page) < near/2 > table) < sentence > ((memory or graphic) <near/2> (window or space or region or area))] Found 17 of 107,305 searched. → Rerun within the Portal

Search within Results

Sort by:

> Advanced Search : > Search Help/Tips

Title Publication

Results 1 - 17 of 17 short listing

1 An interactive design automation system

100%

Stephen Y. H. Su

Proceedings of the tenth design automation workshop on Design automation June 1973

Publication Date

Score

Binder

An interactive design automation system is presented which, after complete implementation, will allow the designer to check the determinacy and dead locks of the system before implementation. The design can be evaluated at various levels and modified interactively. The designer enters his design specification using either graphical representation or design language statements. The translator accepts the input and produces a data base for both the simulator and the logic synthesizer. The syn ...

An efficient virtual machine implementation

100%

Ronald J. Srodawa, Lee A. Bates

Proceedings of the workshop on virtual computer systems March 1973

This paper describes the techniques used to implement an efficient virtual machine facility within MTS for the IBM System/360 Model 67. The goals of the project were to support the IBM Operating System, including the Indexed Sequential Access Method and Teleprocessing capabilities, as a subsystem under MTS with a maximum teleprocessing degradation of 30% for OS/360 programs and complete protection between OS/360 and MTS. The first attempt, using channel program relocation



similar to that em ...

Mirage: a coherent distributed shared memory design

100%

B. Fleisch, G. Popek

ACM SIGOPS Operating Systems Review , Proceedings of the twelfth ACM symposium on Operating systems principles November 1989 Volume 23 Issue 5

Shared memory is an effective and efficient paradigm for interprocess communication. We are concerned with software that makes use of shared memory in a single site system and its extension to a multimachine environment. Here we describe the design of a distributed shared memory (DSM) system called Mirage developed at UCLA. Mirage provides a form of network transparency to make network boundaries invisible for shared memory and is upward compatible with an existing interfac ...

4 The VMP multiprocessor: initial experience, refinements, and

100%

d performance evaluation

D. R. Cheriton , A. Gupta , P. D. Boyle , H. A. Goosen ACM SIGARCH Computer Architecture News, Proceedings of the 15th Annual International Symposium on Computer architecture May 1988

Volume 16 Issue 2

VMP is an experimental multiprocessor being developed at Stanford University, suitable for high-performance workstations and server machines. Its primary novelty lies in the use of software management of the per-processor caches and the design decisions in the cache and bus that make this approach feasible. The design and some uniprocessor trace-driven simulations indicating its performance have been reported previously. In this paper, we present our initial experience with the V ...

Minimizing wasted space in partitioned segmentation Erol Gelenbe , J. C. A. Boekhorst , J. L. W. Kessels Communications of the ACM June 1973 Volume 16 Issue 6

100%

A paged virtual memory system using a finite number of page sizes is considered. Two algorithms for assigning pages to segments are discussed. Both of these algorithms are simple to implement. The problem of choosing the page sizes to minimize the expected value of total wasted space in internal fragmentation and in a page table, per segment, is then solved for a probability density function of segment size which may be expressed as a convex combination of Erlang densities.



100%

Peter J. Denning
ACM Computing Surveys (CSUR) December 1976
Volume 8 Issue 4

100%

Practicing JUDO: Java under dynamic optimizations

Micha? Cierniak , Guei-Yuan Lueh , James M. Stichnoth ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN '00 conference on Programming language design and implementation May 2000

Volume 35 Issue 5

A high-performance implementation of a Java Virtual Machine (JVM) consists of efficient implementation of Just-In-Time (JIT) compilation, exception handling, synchronization mechanism, and garbage collection (GC). These components are tightly coupled to achieve high performance. In this paper, we present some static anddynamic techniques implemented in the JIT compilation and exception handling of the Microprocessor Research Lab Virtual Machine (MRL VM), ...

8 The TLB slice—a low-cost high-speed address translation 100% mechanism

George Taylor , Peter Davies , Michael Farmwald ACM SIGARCH Computer Architecture News , Proceedings of the 17th annual international symposium on Computer Architecture May 1990

Volume 18 Issue 3

The MIPS R6000 microprocessor relies on a new type of translation lookaside buffer — called a TLB slice — which is less than one-tenth the size of a conventional TLB and as fast as one multiplexer delay, yet has a high enough hit rate to be practical. The fast translation makes it possible to use a physical cache without adding a translation stage to the processor's pipeline. The small size makes it possible to include address translation on-chip, even in a tech ...

9 Efficient transparent application recovery in client-server

100%

d information systems

David Lomet, Gerhard Weikum

ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data June 1998 Volume 27 Issue 2

Database systems recover persistent data, providing high database availability. However, database applications, typically residing on client or " middle-tier" application-server





machines, may lose work because of a server failure. This prevents the masking of server failures from the human user and substantially degrades application availability. This paper aims to enable high application availability with an integrated method for database server recovery and tra ...

10 Operating system support for improving data locality on

100%

d CC-NUMA compute servers

Ben Verghese , Scott Devine , Anoop Gupta , Mendel Rosenblum Proceedings of the seventh international conference on Architectural support for programming languages and operating systems September 1996

The dominant architecture for the next generation of shared-memory multiprocessors is CC-NUMA (cache-coherent non-uniform memory architecture). These machines are attractive as compute servers because they provide transparent access to local and remote memory. However, the access latency to remote memory is 3 to 5 times the latency to local memory. CC-NOW machines provide the benefits of cache coherence to networks of workstations, at the cost of even higher remote access latency. Given the larg ...

11 Exokernel: an operating system architecture for

100%

d application-level resource management

D. R. Engler, M. F. Kaashoek, J. O'Toole

ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles December 1995

Volume 29 Issue 5

12 Virtual address translation for wide-address architectures

100%

d Ing-Jye Shyu, Shiuh-Pyng Shieh

ACM SIGOPS Operating Systems Review October 1995 Volume 29 Issue 4

Operating systems employ virtual memory mechanism to provide large address space for programs. The efficiency of the virtual address translation plays an important role in determining system performance. In conventional virtual memory management systems, both the forward-mapped page table scheme and inverted page table scheme are widely used to organize the page tables that record translation data. These two schemes work well for 32-bit architectures, but not for wide address (64-bit) architectu...



3 Joseph D. Touch

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication October 1995 Volume 25 Issue 4

MD5 is an authentication algorithm proposed as the required implementation of the authentication option in IPv6. This paper presents an analysis of the speed at which MD5 can be implemented in software and hardware, and discusses whether its use interferes with high bandwidth networking. The analysis indicates that MD5 software currently runs at 85 Mbps on a 190 Mhz RISC architecture, a rate that cannot be improved more than 20-40%. Because MD5 processes the entire body of a packet, this data ra ...

14 Efficient software-based fault isolation

100%

Robert Wahbe , Steven Lucco , Thomas E. Anderson , Susan L. Graham

ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles December 1993
Volume 27 Issue 5

15 Processor allocation for a class of hypercube-like

100%

d supercomputers

N. G. Haravu , S. G. Ziavras

Proceedings of the 1992 ACM/IEEE conference on Supercomputing December 1992

16 A theory for memory-based learning

100%

Jyh-Han Lin , Jeffrey Scott Vitter

Proceedings of the fifth annual workshop on Computational learning theory July 1992

A memory-based learning system is an extended memory management system that decomposes the input space either statically or dynamically into subregions for the purpose of storing and retrieving functional information. The main generalization techniques employed by memory-based learning systems are the nearest-neighbor search, space decomposition techniques, and clustering. Research on memory-based learning is still in its early stage. In particular, there are very few rigorous theoretical r ...

17 The interaction of architecture and operating system design

100%

Thomas E. Anderson, Henry M. Levy, Brian N. Bershad, Edward D. Lazowska





ACM SIGARCH Computer Architecture News , Proceedings of the fourth international conference on Architectural support for programming languages and operating systems April 1991 Volume 19 Issue 2

Results 1 - 17 of 17 short listing

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.